

Year	1970	1971	1972	1973	1974	1975	1976	1977	1978	1979	1980	1981	1982	1983	1984	1985	1986	1987	1988	1989	1990	1991	1992	1993	1994	1995	1996	1997	1998	1999	2000	2001	2002	2003	2004	2005	2006	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	2026	2027	2028	2029	2030	2031	2032	2033	2034	2035	2036	2037	2038	2039	2040	2041	2042	2043	2044	2045	2046	2047	2048	2049	2050	2051	2052	2053	2054	2055	2056	2057	2058	2059	2060	2061	2062	2063	2064	2065	2066	2067	2068	2069	2070	2071	2072	2073	2074	2075	2076	2077	2078	2079	2080	2081	2082	2083	2084	2085	2086	2087	2088	2089	2090	2091	2092	2093	2094	2095	2096	2097	2098	2099	2100
1970	1971	1972	1973	1974	1975	1976	1977	1978	1979	1980	1981	1982	1983	1984	1985	1986	1987	1988	1989	1990	1991	1992	1993	1994	1995	1996	1997	1998	1999	2000	2001	2002	2003	2004	2005	2006	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	2026	2027	2028	2029	2030	2031	2032	2033	2034	2035	2036	2037	2038	2039	2040	2041	2042	2043	2044	2045	2046	2047	2048	2049	2050	2051	2052	2053	2054	2055	2056	2057	2058	2059	2060	2061	2062	2063	2064	2065	2066	2067	2068	2069	2070	2071	2072	2073	2074	2075	2076	2077	2078	2079	2080	2081	2082	2083	2084	2085	2086	2087	2088	2089	2090	2091	2092	2093	2094	2095	2096	2097	2098	2099	2100	

Both citizens of the United States of America

ANALYZING MOTION OF CHARACTERISTICS IN IMAGES

BACKGROUND

There are many ways in which motion may be estimated between two images. This motion may be described by a set of motion parameters that describe motion of luminance of pixels from a first image to a second image. These motion parameters may be defined at a time associated with either or both of the first and second images, or may be defined at a time between the first and second images. Thus, a vector for each pixel describes the motion of the luminance of the pixel from one image to the next. Motion also may be described by a parameterized motion model, which may be translational, using two parameters, affine, using six parameters, or projective, using eight parameters, and that is defined for a region of an image, or an entire image. An estimate of a single parameterized motion model for a user-defined region of an image is useful for stabilization and tracking applications. An estimate of translational motion for every pixel in the image may be used for sample rate conversion and morphing applications. This motion estimate may be computed by using a gradient-based method, of which an example is a technique referred to as computing the "optical flow" between the images, or by using a correlation-based method.

Such motion parameters may be estimated by relying on what is known as a constant brightness constraint. The assumption is that the total luminance from one image to the next is constant. Two images, for example in RGB format, are converted from the existing format to a single luminance component, typically the luminance component of a YCrCb format image. Parameters are first estimated on a reduced resolution image, then propagated to a higher resolution version of the image. Details about implementations of such motion analysis may be found in several references, including, but not limited to "Hierarchical Model-Based Motion Estimation," by J.R. Bergen et al., in Proceedings of Second European Conference on Computer Vision, pages 237-252, Springer-Verlag, 1992; and "Hierarchical Model-Based Frame Rate Conversion," by J.R. Bergen et al, Technical Report, David Sarnoff Research Center, 1990; and "The Computation of Optical Flow, by S.S. Beauchemin and J.L. Barron,

ACM Computing Surveys, Vol. 27, No. 3, September 1995, pp. 433-467, which are hereby incorporated by reference.

SUMMARY

In some neighboring images in a video or film sequence, the constant brightness constraint does not hold. Such a condition may arise because of a change in an object's position relative to light sources, an object's specularity, an overall luminance change, or a lack of similarity between the images. In calculating motion between two images, instead of generating an image of a single component comprised of the luminance component of an image, a single channel image may be generated from the image based on some other desired characteristic. Given a desired characteristic (such as edge strength or edge magnitude) in an image, a function measures the strength of the desired characteristic in a region around a pixel in an image. A range of values can represent the likelihood, or measure of confidence, of the occurrence of the desired characteristic in the region around the pixel. Thus, each pixel in the single channel image has a value from the range of values that is determined according to a function. This function operates on a neighborhood in the input image that corresponds to the pixel in the single channel image, and measures the likelihood of occurrence of, or strength of, the desired characteristic in that neighborhood. Two single channel images generated from two images are analyzed to provide a motion estimate that indicates how the location of characteristics in the images changes from one image to the next image. If the desired characteristic is an edge magnitude or edge strength, then the motion is effectively estimated using a constant edge constraint.

Accordingly, in an aspect, motion analysis is performed on two images by generating a single channel image for each of the two input images according to a function that measures, for each pixel, occurrence of a desired characteristic, other than luminance alone, in the input images at each pixel location to provide a value for an output pixel in the single channel image from a range of values. An estimate of motion of the desired characteristic between the two images is computed using the single channel images generated for the two input images. The input images may be processed according to the estimate of motion. The desired characteristic may be edge magnitude.

The estimate of motion may be used to process the input images to generate several images from the first image to the second image.

In another aspect, image processing is performed on two images by computing an estimate of motion between the two images according to a constant edge constraint. The images may be processed according to the estimate of motion.

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is a dataflow diagram of a system that analyzes motion of characteristics in images.

Fig. 2 is a dataflow diagram illustrating more detail of an example computation of a motion estimate.

Fig. 3 is a dataflow diagram of a system that analyzes motion based on a constant edge constraint.

Fig. 4 is a graph of a function that performs linear post-processing of edge magnitude.

Fig. 5 is a graph of a function that performs nonlinear post-processing of edge magnitude.

DETAILED DESCRIPTION

Fig. 1 is a dataflow diagram of a system that analyzes motion of characteristics in images. In Fig. 1, each image 100, 102 is processed by characteristic measurement processors 104, 106 respectively, to produce a single channel image 108, 110 based on a desired characteristic of that image. Although two characteristic measurement processors are shown, the images could be processed serially by one characteristic measurement processor.

The characteristic measurement processor implements a function, examples of which are provided below, that measures the occurrence of a desired characteristic in a region around a pixel in an image. A range of values can represent the likelihood, or measure of confidence, of the occurrence of the desired characteristic in the region around the pixel. Thus, each pixel in the single channel image has a value from the range of values that is determined according to a function. This function operates on a neighborhood in the input image that corresponds to the pixel in the single channel image, and measures the likelihood of occurrence of, or strength of, the desired characteristic in that neighborhood.

Two single channel images 108, 110 generated from the two input images 100, 102 are analyzed by a motion estimation process 112 to provide a motion estimate 114 that indicates how the location of a characteristic changes from one image 100 to the next image 102. The motion estimation process 112 also may use the original input images as indicated at 118, in a manner described below in connection with Fig. 2. The motion estimation may be based on optical flow, such as described in the references noted above, or other equivalent motion estimation process. Such computations generally are gradient-based or correlation-based. The motion estimate may be in the form of a per-pixel vector map, or any conventional motion model, such as a model for a region of an image or a model for the entire image.

After the motion estimate 114 is computed, the motion estimate 114 may be used to perform post processing motion operations 116, such as warping, morphing, motion blurring, stabilization, image sharpening, mosaic generation or other effects, on the input images 100, 102. Various post-processing operations, and methods for computing optical flow, are described in a related application serial number 09/657,699, filed September 8, 2000, and U.S. patent applications entitled, "Correcting Motion Vector Maps for Image Processing" by Katherine Cornog and Randy Fayan and "Interpolation of a Sequence of Images Using Motion Analysis" by Katherine Cornog, et al., filed on even date herewith, and hereby incorporated by reference.

Fig. 2 is a dataflow diagram illustrating more detail of an example computation of a motion estimate that may use both the input images 100, 102 (Fig. 1) and the single channel images 108, 110 (Fig. 1). The input images 200, 202 are processed by luminance selection 204, 206 to generate a luminance or grey-scale image 208, 210. The luminance images 208, 210 are blended with their respective single channel images (based on characteristic measurement) 212, 214 (see 108, 110 in Fig. 1) by blend operations 216, 218 to produce output images 220, 222. The blend operations may be implemented, for example, by a typical alpha blend of the function $C_{x,y} = \alpha A_{x,y} + (1-\alpha)B_{x,y}$, where $C_{x,y}$ is a pixel at coordinates (x,y) in the output image, α is a blend value in the range of zero to one, $A_{x,y}$ is a pixel in one of the input images, and $B_{x,y}$ is a pixel in the other of the input images. The blend value α (226, 228) may be specified by a user through any conventional user interface technique for obtaining user input. If $\alpha=0$ or $\alpha=1$, one of the

images input to the blend is output, and the other image has no contribution to that output.

The blended images 230, 232 are used to compute a motion estimate, as indicated at 234. As noted above, the motion estimation 234 may be based on optical flow, such as described in the references noted above, or other equivalent motion estimation. The motion estimate may be in the form of a per-pixel vector map or any conventional motion model, such as a parameterized model for a region of the image or for the entire image.

Examples of the kinds of characteristics that may be measured by characteristic processors 104, 106 in Fig. 1 will now be described. Characteristics for which it might be desirable to estimate motion include characteristics such as edge magnitude, proximity to a selected color, or other characteristics. In general, a function that describes the likelihood of occurrence or measure of the strength of the characteristic, over a range of values, is used to generate a single channel image. The range of values may be, for example, a discrete integer range (e.g., 0 to 255) or a range of fractional values expressed in fixed-point or floating-point format.

By measuring edge magnitude, for example, some artifacts generated by using optical flow based only on luminance may be reduced or eliminated. For example, with images that have structurally similar but visually dissimilar features, the results from image processing based on motion estimated using optical flow can be poor. By measuring optical flow based on edge magnitude, a better match of the structures of the two images may be obtained.

An example characteristic measurement processor that measures edge magnitude will now be described. In this example, each of the input images is processed to create an edge magnitude image, indicative of the strength of an edge at each pixel in the image. The edge magnitude is computed by combining the output of differential operators, called horizontal and vertical edge operators. Linear and nonlinear post processing of the edge magnitude images may be applied to adaptively center the edge magnitude in a specified range. The sign of the edge magnitude may be retained as well, for avoiding matching of light to dark transitions with dark to light transitions.

The following formulas describe the edge magnitude operation as it is performed for each pixel in an input image, to obtain values del1 and del2 that represent the edge magnitude. A first embodiment, described using MATLAB notation, is as follows:

```
del1 = sqrt(dx1.*dx1+dy1.*dy1);
```

```
del2 = sqrt(dx2.*dx2+dy2.*dy2);
```

where dx1 is the x derivative of the first image, dy1 is the y derivative of the first image, dx2 is the x derivative of the second image, and dy2 is the y derivative of the second image. Each derivative may be calculated based on a 3-tap filter centered on the pixel, with coefficients of zero for the center tap and 0.5 and -0.5 for the taps for the adjacent pixels. Other derivative filters may be used.

Another method for computation of edge magnitude is as follows:

```
del1 = abs(dx1) + abs(dy1);
```

```
del2 = abs(dx2) + abs(dy2);
```

where dx1 is the x derivative of the first image, dy1 is the y derivative of the first image, dx2 is the x derivative of the second image, and dy2 is the y derivative of the second image.

The edge magnitude for each pixel optionally may be post-processed to normalize it in the range (e.g. 8-bits) of an image. For example, a scale factor may be computed and applied for each pixel as follows:

```
md = mean(del1(:)+del2(:))/2;
```

```
stdd=std(del1(:)+del2(:))/2;
```

```
psi=4;
```

```
scale = 255/(md + psi*stdd);
```

```
offset = 0;
```

```
b1 = scale*del1+offset; and
```

```
b2 = scale*del2+offset,
```

where b1 and b2 are the normalized versions of the edge magnitude images.

After optional scaling, the values are clipped to the range by limiting all values that are less than the minimum (e.g., zero) to the minimum and all values greater than the maximum (e.g., 255) to the maximum. Further post processing may be optionally applied to enhance contrast and reduce the effect of low amplitude edges. The a linear

function may be applied, such as shown in Fig. 4. A nonlinear function also may be used, such as described by the following formula. The values of alpha and beta may vary in these formulas.

$m = \text{mean}(b1(:)+b2(:))/2;$

$\text{sigma} = \text{std}(b1(:)+b2(:))/2;$

$\text{alpha}=1.5;$

$\text{beta}=.1$

$c1=255./(1+\exp(-(b1-\text{alpha}*m)./(\text{beta}*\text{sigma})));$ and

$c2=255./(1+\exp(-(b2-\text{alpha}*m)./(\text{beta}*\text{sigma}))),$

where c1 and c2 are the versions of the edge magnitude images that have been post-processed. Fig. 5 illustrates a graph of such a function.

By processing the input images to determine edge magnitude, and processing optical flow on the edge magnitude of the input images, the effect is that optical flow is computed based on a constant edge constraint as opposed to a constant brightness constraint. Therefore, as shown in Fig. 3, a system is provided that receives input images 300, 302 and determines a motion estimate 304 using a motion estimation process 306 based on a constant edge constraint. The input images 300, 302 may be processed by a post-processing operations 308 using the motion estimate so generated.

Another example of another characteristic that may be measured is proximity to a specified color. Proximity is a measure of distance or similarity of two colors in a color space. A single channel image may be generated from an input image such that a pixel in the input image that is equal to the specified color is converted to a pixel in an output image at one extreme of the range of output values, e.g., a value representative of white. All other pixels in the input image are converted to pixels in the output image that have a value representative of their distance from the specified color.

Having now described an example embodiment, it should be apparent to those skilled in the art that the foregoing is merely illustrative and not limiting, having been presented by way of example only. Numerous modifications and other embodiments are within the scope of one of ordinary skill in the art and are contemplated as falling within the scope of the invention.

What is claimed is: